

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

| | | |
|--------------------------------|---|----------------------------|
| In re Application of: | § | Examiner: Chou, Alan S |
| Bala Dutt | § | Group Art Unit: 2151 |
| Ajay Kumar | § | Atty. Dkt. No.: 5681-14900 |
| Venugopal Rao K | § | |
| Sankara R. Bhogi | § | |
| Srinivasan Kannan | § | |
| | § | |
| Serial No. 10/655,346 | § | |
| | § | |
| Filed: September 4, 2003 | § | |
| | § | |
| For: Identity for Data Sources | § | |
| | § | |

APPEAL BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir/Madam:

Further to the Notice of Panel Decision mailed September 2, 2009, Appellants present this Appeal Brief. Appellants respectfully request that the Board of Patent Appeals and Interferences consider this appeal.

I. REAL PARTY IN INTEREST

As evidenced by the assignment recorded at Reel/Frame 104470/0690, the subject application is owned by Sun Microsystems, Incorporated, a corporation organized and existing under and by virtue of the laws of the State of Delaware, and having its principal place of business at 4150 Network Circle, Santa Clara, CA 95054.

II. RELATED APPEALS AND INTERFERENCES

No other appeals, interferences or judicial proceedings are known which would be related to, directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

Claims 1-23 are pending in the application and stand finally rejected. The rejection of claims 1-23 is being appealed. A copy of claims 1-23 is included in the Claims Appendix herein below.

IV. STATUS OF AMENDMENTS

No amendments have been submitted subsequent to the final rejection.

V. SUMMARY OF CLAIMED SUBJECT MATTER

Independent claim 1 is directed to a system that includes an application server. (See, e.g., Fig. 1; and p. 8, line 3.) The system also includes one or more backend systems coupled to the application server. (See, e.g., Fig. 1; and p. 12, line 22.) The one or more backend systems include a plurality of distinct data resources. (See, e.g., Fig. 1; and p. 2, lines 5-6.) The application server has an application configured to initiate requests for connections with the plurality of distinct data resources. (See, e.g., Fig. 2; and p. 10, lines 6-7.) The application server also includes a plurality of data sources configured to provide connections with the plurality of distinct data resources. (See, e.g., Fig. 2; and p. 4, lines 16-18.) The application server is configured to associate an identity with each of the plurality of data sources (see, e.g., Figs. 2-3; and p. 4, line 4) and to use the identity to determine whether one of the plurality of data sources provides connections to the same data resource as another of the plurality of data sources (see, e.g., Fig. 3; and p. 9, lines 21-24), where each identity is unique to one of the plurality of distinct data resources (see, e.g., Fig. 1; and p. 4, lines 6-7), and where multiple ones of the data sources have the same identity (see, e.g., Figs. 1 and 4; p. 9, lines 21-24; and p. 14, lines 1-2).

Independent claim 12 is directed to a method that includes receiving notification of a request for a connection. (See, e.g., Fig. 3; and p. 9, lines 15-16.) The method also includes ascertaining an identity of a data source associated with the request (see, e.g., Fig. 3; and p. 4, line 8), where the data source is configured to provide the connection to one of a plurality of distinct data resources (see, e.g., Figs. 1 & 2; p. 8, lines 12-15; and p. 9, line 15), and where the identity is unique to the one of the plurality of distinct data resources (see, e.g., Fig. 1; and p. 4, lines 6-7). The method calls for comparing the identity of the data source with respective identities of multiple data sources with existing connections (see, e.g., Fig. 3; and p. 4, lines 8-9), where the identity of each of the multiple data sources is unique to a specific one of the plurality of distinct data resources (see, e.g., Fig. 2; and p. 4, lines 6-7). The method implements providing an existing connection if an identity match is found with one of the data sources with existing

connections (see, e.g., Fig. 3; and p. 4, lines 9-10). The method also performs providing a new connection if no identity match is found (see, e.g., Fig. 3; and p. 4, lines 10-11).

Independent claim 17 is directed to a method that includes receiving a request for a connection with a participant in a transaction. (See, e.g., Fig. 3; and p. 9, lines 5-16.) The method also includes ascertaining an identity of a local data source associated with the request. (See, e.g., Fig. 3; p. 4, line 8; p. 4, lines 28-29; and p. 8, lines 8-10.) The method calls for attempting to identify a data source that is already participating in the transaction whose identity matches the identity of the data source associated with the request, where the attempting comprises comparing the ascertained identity to identities for a plurality of data sources (see, e.g., Fig. 3; p. 9, lines 5-10; and p. 11, lines 22-25), where each of the plurality of data sources is configured to provide a connection with one of a plurality of distinct data resources (see, e.g., Fig. 2; p. 4, lines 16-18; and p. 10, lines 8-9 and 20-21), and where each identity is unique to one of the plurality of distinct data resources (see, e.g., Fig. 1; and p. 4, lines 6-7). The method involves sharing an existing connection associated with the identity if a data source with a matching identity is found (see, e.g., Fig. 3; and p. 4, lines 9-10), and providing a new connection if no data source with a matching identity is found (see, e.g., Fig. 3; and p. 4, lines 10-11).

Independent claim 18 is directed to a computer-accessible medium comprising program instructions (see, e.g., Fig. 8; and p. 18, lines 16-26), where the program instructions are computer-executable to receive notification of a request for a connection (see, e.g., Fig. 3; and p. 9, lines 15-16); ascertain an identity of a data source associated with the request (see, e.g., Fig. 3; and p. 4, line 8); determine whether the identity of a data source from among a plurality of local and global data sources each having an existing connection matches the identity of a data source associated with the request (see, e.g., Fig. 3; p. 4, lines 7-11; and p. 8, lines 13-14), where each of the plurality of local and global data sources is configured to provide a connection with one of a plurality of distinct data resources (see, e.g., Figs. 1& 2; p. 4, lines 16-18; p. 8, lines 12-15; and p. 9, line 15), and where each of the plurality of local and global data sources has an identity unique to one of the plurality of distinct data resources (see, e.g., Fig. 1; and p. 4, lines 6-

7); provide the existing connection if an identity match is found (see, e.g., Fig. 3; and p. 4, lines 9-10); and provide a new connection if no identity match is found (see, e.g., Fig. 3; and p. 4, lines 10-11).

Independent claim 23 is directed to computer-accessible medium comprising program instructions (see, e.g., Fig. 8; and p. 18, lines 16-26), where the program instructions are computer-executable to receive a request for a connection with a participant in a transaction (see, e.g., Fig. 3; and p. 9, lines 5-16); ascertain an identity of a local data source associated with the request (see, e.g., Fig. 3; p. 4, line 8; p. 4, lines 28-29; and p. 8, lines 8-10); attempt to identify a data source that is already participating in the transaction whose identity matches the identity of the data source associated with the request, where to attempt to identify comprises comparing the ascertained identity to identities for a plurality of data sources (see, e.g., Fig. 3; p. 9, lines 5-10; and p. 11, lines 22-25), where each of the plurality of data sources is configured to provide a connection with one of a plurality of distinct data resources (see, e.g., Fig. 2; p. 4, lines 16-18; and p. 10, lines 8-9 and 20-21), and where each identity is unique to one of the plurality of distinct data resources (see, e.g., Fig. 1; and p. 4, lines 6-7); share an existing connection associated with the identity if a data source with matching identity is found (see, e.g., Fig. 3; and p. 4, lines 9-10); and provide a new connection if no data source with matching identity is found (see, e.g., Fig. 3; and p. 4, lines 10-11).

The summary above describes various examples and embodiments of the claimed subject matter; however, the claims are not necessarily limited to any of these examples and embodiments. The claims should be interpreted based on the wording of the respective claims.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 1-5, 8-11, 12, 15-18 and 21-23 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Ng (U.S. Patent 6,411,956) in view of Felt et al. (U.S. Patent 7,080,119) (hereinafter “Felt”).

2. Claims 6, 7, 13, 14, 19 and 20 stand finally rejected under 35 U.S.C. § 103(a) as being unpatentable over Ng in view of Felt and Yousefi’zadch (U.S. Publication 2004/0030739).

VII. ARGUMENT

First Ground of Rejection:

The Office Action rejected claims 1-5, 8-11, 12, 15-18 and 21-23 under 35 U.S.C. § 103(a) as being unpatentable over Ng (U.S. Patent 6,411,956) in view of Felt et al. (U.S. Patent 7,080,119) (hereinafter “Felt”). Appellants respectfully traverse this rejection for at least the following reasons.

Claims 1, 3-5, and 8-11:

1. Ng in view of Felt clearly fails to teach or suggest *an application server that comprises: an application configured to initiate requests for connections with a plurality of distinct data resources; a plurality of data sources configured to provide connections with the plurality of distinct data resources; and wherein the application server is configured to associate an identity with each of the plurality of data sources and to use the identity to determine whether one of the plurality of data sources provides connections to the same data resource as another of the plurality of data sources.*

The Final Office Action refers to Ng at column 4, lines 10-17. The cited lines merely present a flow chart for the operation of the adapter described at column 3. According to Ng, instead of creating a new physical database connection whenever one of the components of the *same* transaction requests a connection to the *same* database, Ng’s adapter looks to see if there is already an existing connection to that *same* database made by some component of that *same* transaction. For this purpose, Ng consults Table 60 that maps global transaction identifiers to the physical database connections to database 601 [Figure 6]. If an existing connection to the *same* database 601 for the *same* transaction is found, the requesting component is directed to use the existing connection. In this manner, all the work of the *same* transaction that is associated with the *same* database 601 can be coordinated under the *same* transaction [column 3, lines 40-49]. The

purpose of Ng's adapter is to provide a limited version of JDBC 2.0 distributed transaction support for databases that don't have a JDBC 2.0 driver [column 1, lines 48-50]. Ng pertains to a system using only existing JDBC 1.0 drivers where for a given transaction there is only *one* physical connection to *one* database [column 3, lines 30-34].

As noted above, Ng's Table 60 maps global transaction identifiers to the physical database connections to database 601 [Figure 6] so that if an existing connection to the *same* database 601 for the *same* transaction is found, the requesting component is directed to use the existing connection. Ng clearly does not teach *a particular application that is configured to initiate requests for connections with a plurality of distinct data resources* (e.g., multiple distinct databases). In fact, the Final Office Action acknowledges this fact at paragraph 5 on page 3: "Ng does not disclose expressly the use of distinct data resources."

The Final Office Action refers to Ng's Table 60 and the associated first paragraph of column 4. As explained above, Ng's Table 60 lists global *transaction identifiers* and their physical database connections to a *single* database 601. Ng's global transaction identifiers are not *data sources configured to provide connections with the plurality of distinct data resources*. They merely identify various transactions whose components are all part of a particular process [column 4, lines 10-32]. Ng's table is just a table, not a plurality of data sources in that application server that perform the function of providing connections with the plurality of distinct data resources. The components of a transaction *request* connections to a data resource. Transactions and their components are not *data sources configured to provide connections to a plurality of distinct data resources*. The physical database connections are not *data sources configured to provide connections to distinct data resources*, either. The physical database connections *are* the actual connections themselves to data resources; they are not *data sources configured to provide connections to data resources*. Nothing in Table 60 or elsewhere in Ng teaches *a plurality of data sources configured to provide connections with the plurality of distinct data resources*.

The Final Office Action refers to Ng at column 4, lines 15-27. The cited lines refer again to Table 60, and to a component of a transaction requesting a connection with a database. As explained above, neither the transactions nor the physical database connections to which Table 60 provides a mapping are data sources configured to provide connections with the plurality of distinct data resources. Neither do the cited lines 15-27 of column 4 indicate associating an identity with data sources that are configured to provide connections with the plurality of distinct data resources. Instead, when a component of a transaction requests a connection to database 601, Ng simply checks Table 60 to determine whether there is an existing physical connection to database 601 that is already associated with the component's transaction. If so, the component will use the existing physical connection. The table in Ng associates a *transaction identifier* with an actual connection, not with the data source that provides the connections. Also, Ng does not use his table determine whether one of the multiple data sources already functions to provide connections to the same data resource as another of the multiple data sources does. Nor does Felt, whether considered alone or in combination with Ng, teach these aspects of Appellants' claim 1.

For at least these above reasons, a *prima facie* rejection has not been established.

2. Ng in view of Felt clearly fails to teach or suggest wherein each identity is unique to one of the plurality of distinct data resources, and wherein multiple ones of the data sources have the same identity.

Ng does not teach that the data source identities are each unique to one of the plurality of distinct data resources, i.e., that each *data source identity* corresponds to *exactly one* of the plurality of distinct data resources. Neither does Ng teach that multiple ones of the data sources have the same identity, so that, as a logical consequence, multiple ones of the *data sources* are associated by their identities with the very same single *data resource* out of the plurality of distinct data resources, and with no other data resource. The Examiner seems to rely on Felt to remedy these deficiencies in the teaching of Ng. On page 3 of the Final Action, the Examiner admits that Ng does not

disclose “the use of distinct resources and use of an identifier to delegate the task to a distinct data sources,” citing a paragraph in Felt. However, claim 1 does not recite “the use of distinct resources and use of an identifier to delegate the task to a distinct data sources.” Moreover, the cited paragraph of Felt beginning at line 60 of column 8 describes a plurality of servers which receive a client transaction, and a server to handle transaction commit processing and communicate the result of the transaction commit process to the client process. The Examiner makes no attempt to explain how this teaching of Felt relates to the limitations recited in Appellants’ claim 1. There is nothing in the cited text that even remotely suggests that the data source identities are each unique to exactly one of the plurality of distinct data resources, i.e., that each data source identity corresponds to exactly one of the plurality of distinct data resources, nor that multiple ones of the data sources have the same identity, so that, as a logical consequence, multiple ones of the data sources are associated by their identities with the very same single data resource out of the plurality of distinct data resources, and with no other data resource.

The Advisory Action refers to Felt’s “Background of the Invention” section, citing text in column 6. The cited text contrasts distributed transactions with local transactions, and describes participants in distributed transactions, including a transaction manager. The Advisory Action makes the completely unsupported and irrelevant assertion that the working of a standard transaction manager as described in Felt’s background section “has the same function of determining unique data as claimed.” The Examiner completely ignores the actual wording of the claim. Nothing in Felt even remotely suggests that the data source identities are each unique to exactly one of the plurality of distinct data resources, i.e., that each data source identity corresponds to exactly one of the plurality of distinct data resources, nor that multiple ones of the data sources have the same identity, so that, as a logical consequence, multiple ones of the data sources are associated by their identities with the very same single data resource out of the plurality of distinct data resources, and with no other data resource, as recited in claim 1. In fact, neither reference teaches these aspects of Appellants’ claim 1.

Therefore, whether taken individually or in combination, Ng and Felt do not teach the limitations of Appellants' claim 1.

For at least these above reasons, a *prima facie* rejection has not been established.

3. The Examiner has failed to provide a proper reason to combine the references.

The intended purpose of Ng pertains to a system that does *not* have a JDBC 2.0 Standard Extension implementation, and further requires that all the components of a transaction must be part of the same process in order for the method to work [column 3, lines 30-34; column 4, lines 28-33]. Ng's teachings are specific to systems *not supporting* JDBC 2.0 with its Standard Extension API support for distributed transactions using the standard two-phase commit protocol. Therefore, Ng's teachings purposefully do not apply to allowing multiple JDBC connections to multiple databases under the same global transaction. The Final Office Action proposes incorporating Felt's "distinct data resources" into Ng, alluding to two-phase commit protocol for distributed transactions that is supported by the JDBC 2.0 Standard Extension API. However, the whole point of Ng's teachings is to take requests for multiple connections with the *same* database and redirect them to a *single* physical connection to that *same* database, where the requests are all made by components of the same transaction, thus skirting the normal JDBC 1.0 limitations in this regard. There would be no point in doing this if support allowing multiple connections to multiple databases under the same global transaction were available, as it is when JDBC 2.0 Standard Extension is available. Imposing Felt's architecture upon Ng would render Ng's method unsatisfactory for its intended purpose. Thus, there is no suggestion or motivation to make the proposed modification of Ng. *In re Gordon*, 733 F.2d 900, 221 USPQ 1125 (Fed. Cir. 1984).

In a system that supports true global transactions involving multiple distinct data resources (such as a system implementing the JDBC 2.0 Standard Extension), there would be no need to employ Ng's table that maps transactions to connection, which is the

heart of the Examiner's rejection. Thus, the proposed modification would explicitly not result in Appellants' claimed invention. Therefore, a *prima facie* rejection has not been established.

Claim 2:

Ng in view of Felt clearly fails to teach or suggest wherein in response to the application requesting a connection from one of the plurality of data sources, a data source ID manager is configured to ascertain the identity of the data source from which the connection was requested and determine whether the identity matches the identity of any other of the plurality of data sources.

The Final Office Action cites Ng's adapter 523 and text found in the second paragraph of column 4, asserting that "adapter 523 determine connection association" without further elaboration. Once again the Examiner has completely ignored the wording of the claim. Claim 2 does not recite "determine connection association." As explained above, cited Table 60 maps global transaction identifiers to the physical database connections with database 601. In Ng, there is no teaching of a data source ID manager for ascertaining the identity of the data source from which the connection was requested. Appellants' data source is an entity configured to *provide connections* with the plurality of distinct data resources. Appellants' data *source* is not the same as Appellants' data *resource*. Ng does not compare identities of various data *sources* as recited in claim 2. When Ng's component of a transaction requests a connection to the same database 601, Ng simply checks Table 60 to determine whether there is an existing physical connection to same database 601 that is already associated with the component's transaction. There is no data source ID manager configured to ascertain the identity of the data source from which the connection was requested and determine whether the identity matches the identity of any other of the plurality of data sources.

Claims 12, 15, 16, and 18:

Ng in view of Felt clearly fails to teach or suggest *ascertaining an identity of a data source associated with a request for a connection, wherein the data source is configured to provide the connection to one of a plurality of distinct data resources, and wherein said identity is unique to said one of said plurality of distinct data resources, and comparing said identity of the data source requested to provide the connection with respective identities of multiple data sources with existing connections, wherein the identity of each of the multiple data sources is unique to a specific one of said plurality of distinct data resources*, as recited in claim 12.

The Examiner has not fully addressed the limitations recited in Appellants' claims 12 and 18. The Examiner rejected those claims on the same basis as claim 1. However, claim 12 and 18 are not worded the same as claim 1. Therefore, a *prima facie* rejection has not been stated for claims 12 and 18.

Appellants' arguments presented above regarding claim 1 prove that Ng in view of Felt does not teach or suggest *multiple data sources providing connections to one of a plurality of distinct data resources, nor identities for the multiple data sources which are unique to a specific one of the plurality of distinct data resources*. Moreover, as explained in the arguments regarding claim 1, Ng's Table 60 maps *transaction* identifiers to *physical database connections* to a single database 601, which has nothing to do with *ascertaining the identity of a data source associated with a request for a connection to a data resource, where the data source is configured to provide a connection to one of a plurality of distinct data resources, and comparing the ascertained identity of the data source with respective identities of multiple data sources with existing connections*. Nor does Felt have any relevance to these aspects of Appellants' claimed invention. Thus, the combination of Ng and Felt clearly does not teach or suggest Applicants' invention as recited. Therefore, a *prima facie* rejection has not been established.

Claims 17 and 21-23:

Ng in view of Felt clearly fails to teach or suggest receiving a request for a connection with a participant in a transaction; ascertaining an identity of a local data source associated with the request; attempting to identify a data source that is already participating in the transaction whose identity matches the identity of the data source associated with the request, wherein said attempting comprises comparing the ascertained identity to identities for a plurality of data sources, wherein each of the plurality of data sources is configured to provide a connection with one of a plurality of distinct data resources, and wherein each identity is unique to one of the plurality of distinct data resources; sharing an existing connection associated with the identity of a data source with a matching identity is found; and providing a new connection if no data source with a matching identity is found.

The Examiner has never fully addressed the above limitations. Instead, the Examiner rejected claims 17 and 23 on the same basis as claim 1. Claims 17 and 23 are not worded in the same way as claim 1. **Therefore, no *prima facie* rejection has been stated for these claims.**

Substantial portions of the arguments presented above regarding claim 1 also apply to claims 17 and 23. For example, the arguments presented above in regard to claim 1 prove that Ng in view of Felt does not teach or suggest *a plurality of data sources configured to provide a connection with one of a plurality of distinct data resources, nor identities for the plurality of data sources which are unique to a specific one of the plurality of distinct data resources*. Moreover, as explained in the arguments regarding claim 1, Ng's Table 60 maps *transaction* identifiers to *physical database connections* to a single database 601, which has nothing to do with ascertaining an identity of a local data source associated with a request for a connection with a participant in a transaction, and attempting to identify a data source that is already participating in the transaction whose identity matches the identity of the data source associated with the request, where each of the data sources is configured to provide a connection with one of a plurality of distinct

data resources, and comparing the ascertained identity of the local data source to identities for a plurality of data sources. Nor does Felt have any relevance to these aspects of Appellants' claimed invention. Thus, the combination of Ng and Felt clearly does not teach or suggest Applicants' invention as recited. Accordingly, a *prima facie* rejection has not been established.

Second ground of rejection:

The Office Action rejected claims 6, 7, 13, 14, 19 and 20 U.S.C. § 103(a) as being unpatentable over Ng in view of Felt and Yousefi'zadeh (U.S. Publication 2004/0030739). Appellants respectfully traverse this rejection for at least the reasons given above in regard to the corresponding independent claims.

CONCLUSION

For the foregoing reasons, it is submitted that the Examiner's rejection of claims 1-23 was erroneous, and reversal of the Examiner's decision is respectfully requested.

The Commissioner is authorized any fees that may be due to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5681-14900/RCK.

Respectfully submitted,

/Robert C. Kowert/
Robert C. Kowert, Reg. #39,255
Attorney for Appellants

Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
(512) 853-8850

Date: October 2, 2009

VIII. CLAIMS APPENDIX

The claims on appeal are as follows.

1. A system, comprising:

an application server; and

one or more of backend systems coupled to the application server;

wherein the one or more backend systems comprises a plurality of distinct data resources;

wherein the application server comprises:

an application configured to initiate requests for connections with the plurality of distinct data resources;

a plurality of data sources configured to provide connections with the plurality of distinct data resources; and

wherein the application server is configured to associate an identity with each of the plurality of data sources and to use the identity to determine whether one of the plurality of data sources provides connections to the same data resource as another of the plurality of data sources, wherein each identity is unique to one of the plurality of distinct data resources, and wherein multiple ones of the data sources have the same identity.

2. The system as recited in claim 1, wherein in response to the application requesting a connection from one of the plurality of data sources, a data source ID

manager is configured to ascertain the identity of the data source from which the connection was requested and determine whether the identity matches the identity of any other of the plurality of data sources.

3. The system as recited in claim 2, wherein the data source ID manager is further configured to determine whether any of the data sources with matching identities previously supplied a connection to the application and, if a connection was previously supplied, to return the previously supplied connection to the application.

4. The system as recited in claim 3, wherein if no connection from a data source with a matching identity exists, the data source is configured to forward the request to a corresponding data source to obtain a new connection.

5. The system as recited in claim 4, wherein the connection is a local connection.

6. The system as recited in claim 1, wherein the identity comprises values for one or more data source properties.

7. The system as recited in claim 6, wherein the identity comprises database name, URL, and user name.

8. The system as recited in claim 1, wherein one or more of the plurality of data sources is a data source proxy, wherein the application server is configured to instantiate the data source proxy for an abstract name of a data resource used by an application; ascertain an identity for a true data source; and use the identity to link the proxy to the true data source.

9. The system as recited in claim 8, wherein multiple data source proxies correspond to the same data source identity, wherein in response to the application requesting connections with a same data resource from multiple data source proxies, the

data source proxies from which the connections were requested are configured to forward the connection requests to the data source whose identity corresponds to said proxies.

10. The system as recited in claim 1, wherein in response to a request to instantiate a data source corresponding to an abstract name, the application server is configured to determine an identity for the proposed data source, determine whether any existing data source has a matching identity, instantiate the proposed data source only if no existing data source with matching identity is found.

11. The system as recited in claim 1, wherein the application server further comprises a transaction manager;

wherein in response to a request to commit a transaction the transaction manager is configured to identify a number of data resources participating in the transaction according to connections supplied for unique data source identities;

wherein if the number of data resources participating in the transaction is two or more the transaction manager is configured to commit the transaction utilizing a two-phase commit protocol; and

wherein if only one data resource participating in the transaction the transaction manager is configured to commit the transaction utilizing a one-phase commit optimization.

12. A method, comprising:

receiving notification of a request for a connection;

ascertaining an identity of a data source associated with the request, wherein the data source is configured to provide the connection to one of a plurality of

distinct data resources, and wherein said identity is unique to said one of said plurality of distinct data resources;

comparing said identity with respective identities of multiple data sources with existing connections, wherein the identity of each of the multiple data sources is unique to a specific one of said plurality of distinct data resources;

providing an existing connection if an identity match is found with one of the data sources with existing connections; and

providing a new connection if no identity match is found.

13. The method as recited in claim 12, wherein the identity comprises values for one or more data source property.

14. The method as recited in claim 13, wherein the identity comprises database name, URL, and user name.

15. The method of claim 12, further comprising:

creating one or more data source proxies, wherein multiple data source proxies correspond to the same data source identity;

each data source proxy associated with the same data source identity forwarding connection requests to the same data source whose identity corresponds to said proxies.

16. The method of claim 12, further comprising:

in response to a request to commit a transaction, identifying the number of data resources participating in the transaction according to the number of unique data source identities used to establish data resource connections; and

if the number of data resources participating in the transaction is two or more, committing the transaction utilizing a two-phase commit protocol, and if only one data resource is participating in the transaction, committing the transaction utilizing a one-phase commit optimization.

17. A method, comprising:

receiving a request for a connection with a participant in a transaction;

ascertaining an identity of a local data source associated with the request;

attempting to identify a data source that is already participating in the transaction whose identity matches the identity of the data source associated with the request, wherein said attempting comprises comparing the ascertained identity to identities for a plurality of data sources, wherein each of the plurality of data sources is configured to provide a connection with one of a plurality of distinct data resources, and wherein each identity is unique to one of the plurality of distinct data resources;

sharing an existing connection associated with the identity if a data source with a matching identity is found; and

providing a new connection if no data source with a matching identity is found.

18. A computer accessible medium comprising program instructions, wherein the program instructions are computer-executable to:

receive notification of a request for a connection;

ascertain an identity of a data source associated with the request;

determine whether the identity of a data source from among a plurality of local and global data sources each having an existing connection matches the identity of a data source associated with the request, wherein each of the plurality of local and global data sources is configured to provide a connection with one of a plurality of distinct data resources, and wherein each of the plurality of local and global data sources has an identity unique to one of the plurality of distinct data resources;

provide the existing connection if an identity match is found; and

provide a new connection if no identity match is found.

19. The computer accessible medium as recited in claim 18, wherein the identity comprises values for one or more data source properties.

20. The computer accessible medium as recited in claim 19, wherein the identity comprises database name, URL, and user name.

21. The computer accessible medium as recited in claim 18, wherein the program instructions are further computer-executable to:

create one or more data source proxies, wherein multiple data source proxies correspond to the same data source identity;

wherein each data source proxy associated with the same data source identity forwards connection requests to the same data source whose identity corresponds to said proxies.

22. The computer accessible medium as recited in claim 18, wherein the program instructions are further computer-executable to:

in response to a request to commit a transaction, identify the number of data resources participating in the transaction according to the number of unique data source identities used to establish data resource connections; and

if the number of data resources participating in the transaction is two or more, commit the transaction utilizing a two-phase commit protocol, and if only one data resource is participating in the transaction, commit the transaction utilizing a one-phase commit optimization.

23. A computer accessible medium comprising program instructions, wherein the program instructions are computer-executable to:

receive a request for a connection with a participant in a transaction;

ascertain an identity of a local data source associated with the request;

attempt to identify a data source that is already participating in the transaction whose identity matches the identity of the data source associated with the request, wherein to attempt to identify comprises comparing the ascertained identity to identities for a plurality of data sources, wherein each of the plurality of data sources is configured to provide a connection with one of a plurality of distinct data resources, and wherein each identity is unique to one of the plurality of distinct data resources;

share an existing connection associated with the identity if a data source with matching identity is found; and

provide a new connection if no data source with matching identity is found.

IX. EVIDENCE APPENDIX

No evidence submitted under 37 CFR §§ 1.130, 1.131 or 1.132 or otherwise entered by the Examiner is relied upon in this appeal.

X. RELATED PROCEEDINGS APPENDIX

There are no related proceedings.